

Preparation GES-C01 Store - GES-C01 Valid Examcollection



In order to make all customers feel comfortable, our company will promise that we will offer the perfect and considerate service for all customers. If you buy the GES-C01 study materials from our company, you will have the right to enjoy the perfect service. We have employed a lot of online workers to help all customers solve their problem. If you have any questions about the GES-C01 Study Materials, do not hesitate and ask us in your anytime, we are glad to answer your questions and help you use our GES-C01 study materials well.

Snowflake GES-C01 preparation materials will be the good helper for your qualification certification. We are concentrating on providing high-quality authorized GES-C01 study guide all over the world so that you can clear exam one time. As we all know, the preparation process for an exam is very laborious and time-consuming. We had to spare time to do other things to prepare for Snowflake GES-C01 Exam, which delayed a lot of important things.

>> Preparation GES-C01 Store <<

First-grade Preparation GES-C01 Store to Obtain Snowflake Certification

One way to make yourself competitive is to pass the GES-C01 certification exams. Hence, if you need help to get certified, you are in the right place. GuideTorrent offers the most comprehensive and updated braindumps for GES-C01's certifications. To ensure that our products are of the highest quality, we have tapped the services of GES-C01 experts to review and evaluate our GES-C01 certification test materials. In fact, we continuously provide updates to every customer to ensure that our GES-C01 products can

cope with the fast changing trends in GES-C01 certification programs.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q169-Q174):

NEW QUESTION # 169

A financial institution uses Snowflake Cortex Analyst with strict role-based access control (RBAC) on their Snowflake-hosted LLMs. The security team has granted specific 'CORTEX-MODEL-ROLE' application roles to different analyst teams, ensuring they only access approved models. A new requirement arises to enable Azure OpenAI GPT models for Cortex Analyst to leverage a specific feature. An administrator proceeds to execute:

```
USE ROLE ACCOUNTADMIN;
ALTER ACCOUNT SET ENABLE_CORTEX_ANALYST_MODEL_AZURE_OPENAI = TRUE;
```

Which of the following statements accurately describe the implications of this change?

Enabling 'ENABLE_CORTEX_ANALYST_MODEL_AZURE_OPENAI' will allow Cortex Analyst to use Azure OpenAI models, and the existing model-level RBAC policies for Snowflake-hosted models will continue to apply to those models.

When 'ENABLE_CORTEX_ANALYST_MODEL_AZURE_OPENAI' is 'TRUE', Cortex Analyst can use Azure OpenAI models, but model-level RBAC becomes unavailable for *all* models used by Cortex Analyst, including Snowflake-hosted ones.

User prompts and semantic model metadata will be transmitted outside of Snowflake's governance boundary if Azure OpenAI models are used by Cortex Analyst after this parameter is enabled.

This parameter is only relevant for the 'AI_COMPLETE' function and does not affect model access within Cortex Analyst.

The 'ALTER ACCOUNT' command for this parameter can be executed by any role with 'MODIFY ACCOUNT' privileges, not exclusively 'ACCOUNTADMIN'.

- A. Option A
- B. Option C
- C. Option B
- D. Option D
- E. Option E

Answer: B,C

Explanation:

Option B is correct because when is ' TRUE , cortex Analyst can use Azure OpenAI models, but this setting is incompatible with model-level RBAC, meaning RBAC is not available for any models used by Cortex Analyst when this parameter is enabled. Option C is correct because if Azure OpenAI models are opted in for Cortex Analyst, semantic model files (which are metadata) and user prompts will be processed by Microsoft Azure, a third party, thus transmitting them outside Snowflake's governance boundary. Customer data itself is not shared. Option A is incorrect because the parameter is incompatible with model-level RBAC for 'all' models used by Cortex Analyst. Option D is incorrect as the parameter specifically controls the use of Azure OpenAI models within Cortex Analyst. Option E is incorrect because this parameter can only be set by the ' ACCOUNTADMIN' role.

NEW QUESTION # 170

A financial analytics team is using to extract specific financial metrics (e.g., revenue, profit margin) from quarterly reports and requires the output in a strict JSON format for automated ingestion into a data warehouse. They've encountered issues where the LLM sometimes generates malformed JSON or includes extraneous text. Which of the following approaches will help ensure deterministic, schema-compliant JSON outputs and mitigate these 'hallucinations' related to format?

Utilise the `response_format` argument within `AI_COMPLETE` to specify a detailed JSON schema, including `type`, `properties`, and `required` fields.

For OpenAI (GPT) models, ensure that the `additionalProperties` field is set to `false` in every node of the schema and that the `required` field lists all properties.

Include a prompt instruction such as 'Respond in JSON' to improve adherence, especially for complex tasks, even though 'response_format' is provided.

Set the 'temperature' option to 0 when calling 'AI_COMPLETE' to obtain the most consistent and deterministic JSON outputs.

Use the `SNOWFLAKE_CORTEX_CLASSIFY_TEXT` function to validate the JSON structure after generation, then re-prompt if validation fails.

- A. Option C
- B. Option B
- C. Option D
- D. Option A
- E. Option E

Answer: A,B,C,D

Explanation:

Option A is correct because 'AI_COMPLETE Structured Outputs allows supplying a JSON schema via the 'response_format' argument, which defines the structure, data types, and constraints, ensuring responses conform to the schema and reducing post-processing. Option B is correct as it specifically states requirements for OpenAI (GPT) models when using 'response_format', including 'additionalProperties: false' and explicitly listing all properties in the 'required' field. Option C is correct because, for medium-to-complex tasks, adding 'Respond in JSON' to the prompt can further improve JSON adherence accuracy, even when 'response_format' is specified. Option D is correct as setting 'temperature' to 0 is recommended for the most consistent results, which is crucial for deterministic JSON outputs, regardless of task or model. Option E is incorrect because 'CLASSIFY_TEXT is for categorizing text, not for validating JSON structure, making it an inappropriate tool for this task.

NEW QUESTION # 171

A data engineering team is setting up a Retrieval Augmented Generation (RAG) application using Snowflake Cortex Search to provide contextual answers from customer support transcripts. The transcripts are stored in a Snowflake table named SUPPORT_TRANSCRIPTS. Which of the following statements are crucial considerations or accurate facts regarding the initial setup and configuration of the Cortex Search Service for this use case?

- A. Snowflake recommends using a dedicated virtual warehouse of any size, including X-Large or 2X-Large, for each Cortex Search Service to ensure the fastest possible materialization of search indexes during creation and refresh.
- B. The CREATE CORTEX SEARCH SERVICE command requires that CHANGE_TRACKING = TRUE be enabled on the source table, especially if the role creating the service is not the table owner. This ensures that the service can track updates to the base data.
- C. Cortex Search Services currently support replication and cloning, allowing for easy disaster recovery and geographical distribution of the search index.
- D. Cortex Search is designed to get users up and running quickly with a hybrid (vector and keyword) search engine on text data, handling embedding, infrastructure maintenance, and search quality parameter tuning automatically.
- E. Columns specified in the ATTRIBUTES field during service creation are only used for filtering search results and do not need to be present in the source query.

Answer: B,D

Explanation:

Option A is correct because change tracking is required for the Cortex Search Service to monitor updates to the base table, particularly if the service creator is not the table owner. Option B is incorrect; Snowflake recommends using a dedicated warehouse no larger than MEDIUM for each service, as larger warehouses do not necessarily increase performance for index materialization. Option C is incorrect because columns in the ATTRIBUTES field must be included in the source query. Option D is correct as Cortex Search provides a low-latency, high-quality hybrid (vector and keyword) search engine that automatically manages embedding, infrastructure, and search quality tuning. Option E is incorrect because Cortex Search Services currently do not support replication or cloning.

NEW QUESTION # 172

A new ML Engineer, 'data_scientist_role', has been assigned to a project involving custom machine learning models in Snowflake. They need to gain the necessary permissions to perform the following actions related to Snowflake Model Registry and Snowpark Container Services: 1. Log a custom model into a specified schema a. 2. Deploy that model to an existing Snowpark Container Service compute pool. 3. Call the deployed model for inference using SQL. Which of the following SQL commands grant the 'minimal' required privileges to the for these actions, assuming the compute pool and image repository already exist and are appropriately configured?

- A. `GRANT CREATE MODEL ON SCHEMA my_ml_schema TO ROLE data_scientist_role;`
- B. `GRANT READ ON IMAGE REPOSITORY my_model_images TO ROLE data_scientist_role;`
- C. `GRANT CREATE SERVICE ON SCHEMA my_ml_schema TO ROLE data_scientist_role;`
- D. `GRANT CREATE COMPUTE POOL ON SCHEMA my_ml_schema TO ROLE data_scientist_role;`
- E. `GRANT CREATE ON SNOWPARK CONTAINER SERVICE ON SCHEMA my_ml_schema TO ROLE data_scientist_role;`

Answer: A,C

Explanation:

Option A is correct because the 'CREATE MODEL' privilege on the target schema is required to log a new model (which creates a model object) in the Snowflake Model Registry. Option D is correct because deploying a model to a Snowpark Container Service creates a service object within a schema, which requires the 'CREATE SERVICE' privilege on that schema. The role would also implicitly need 'USAGE' on the specified compute pool. Option B is incorrect. While 'USAGE ON DATABASE' is generally needed for accessing objects within a database, it's a broader prerequisite and not specifically a minimal privilege for the direct model registry actions of logging, deploying, and calling the model. Option C is incorrect because 'CREATE COMPUTE POOL' is for creating the compute pool itself, not for deploying a service 'to' an existing one. The role would need 'USAGE' on the existing compute pool, but not the right to create it from scratch for this scenario. Option E is incorrect because 'READ ON IMAGE REPOSITORY' is required for the 'service' to pull the image from the repository, but the question asks for privileges for the to perform the 'actions' of logging, deploying, and calling. While the role might need to manage or verify the image, this isn't a direct privilege for the user's interaction with the deployed model in the same way 'CREATE MODEL' or 'CREATE SERVICE' are.

NEW QUESTION # 173

A Gen AI developer is implementing a Cortex Search Service for a RAG application and needs to configure the text splitting for optimal performance using SNOWFLAKE.CORTEX.SPLIT_TEXT_RECURSIVE_CHARACTER Which of the following statements represent best practices or outcomes when applying text splitting with this function for Cortex Search in a RAG scenario? (Select all that apply)

- A. The function automatically enriches each text chunk with relevant metadata about its original document, such as author and creation date, for enhanced filtering capabilities in Cortex Search.
- B. Optimal text splitting using this function ensures that the number of input tokens precisely equals the number of output tokens for subsequent LLM calls, thereby minimizing compute costs.
- C. Snowflake recommends splitting text into chunks of no more than 512 tokens for best search results in Cortex Search.
- D. Even when using embedding models with larger context windows (e.g., 8000 tokens), a smaller chunk size is typically preferred for improved retrieval and downstream LLM response quality.
- E. Smaller chunk sizes generally lead to higher retrieval precision for a given query in a RAG system

Answer: C,D,E

Explanation:

Options A, B, and C are correct. Snowflake explicitly recommends splitting text in a search column into chunks of no more than 512 tokens for best search results with Cortex Search. Research indicates that smaller chunk sizes typically result in higher retrieval precision for a given query and improved downstream LLM response quality. This practice is recommended even when longer-context embedding models, such as

with an 8000 token context window, are available, because smaller chunks provide more precise retrieval and more relevant context for the LLM. Option D is incorrect; the sources do not mention that SPLIT_TEXT_RECURSIVE_CHARACTER automatically enriches chunks with metadata. This would typically require additional data processing steps. Option E is incorrect; the primary goal of text splitting is to optimize retrieval and LLM response quality, not to balance input and output token counts for cost reasons. While token counts influence costs, the 512-token recommendation is driven by quality considerations.

NEW QUESTION # 174

.....

GuideTorrent has formulated GES-C01 PDF questions for the convenience of Snowflake GES-C01 test takers. This format follows the content of the Snowflake GES-C01 examination. You can read Snowflake GES-C01 Exam Questions without the limitations of time and place. There is also a feature to print out Snowflake GES-C01 exam questions.

GES-C01 Valid Examcollection: <https://www.guidetorrent.com/GES-C01-pdf-free-download.html>

Snowflake Preparation GES-C01 Store Save your time and money, When you will buy GES-C01 preparation material from GuideTorrent you will get the latest one, Being different from the other GES-C01 exam questions in the market, our GES-C01 practice materials have reasonable ruling price and satisfactory results of passing rate up to 98 to 100 percent, Snowflake Preparation GES-C01 Store At the same time, our professional experts are keeping a close eye on the changes of the exam questions and answers.

Controlling Inheritance from the Parent Container, The reasons people GES-C01 exit independent work is more complex and tied to cyclical economic forces and the state of the traditional job market.

Authorized Snowflake GES-C01: Preparation SnowPro® Specialty: Gen AI Certification Exam Store - High Pass-Rate GuideTorrent GES-C01 Valid Examcollection

Save your time and money. When you will buy GES-C01 Preparation material from GuideTorrent you will get the latest one. Being different from the other GES-C01 exam questions in the market, our GES-C01 practice materials have reasonable ruling price and satisfactory results of passing rate up to 98 to 100 percent.

At the same time, our professional experts are keeping a close eye on the changes of the exam questions and answers, GuideTorrent GES-C01 Questions have helped thousands of candidates to achieve their professional dreams.